

REQUIREMENTS ANALYSIS TITLE: NCEP Information Technology Plan

REQUIREMENTS ANALYSIS NUMBER: NA-W-WD-95-01

ORGANIZATION UNIT: National Centers for Environmental Prediction
National Weather Service

TYPE OF SUBMISSION: New

PARTS SUBMITTED: Part 1

DATE OF SUBMISSION: March 1995

CONTACT PERSONS: Arthur Wick, MAR Coordinator
W/NMCx4
301-763-8648
and
William McCracken
Chief, Computer Operations Section
W/NMC431
301-457-5031

Table of Contents

1.	INFORMATION REQUIREMENTS	1
	Current Organization of the National Centers	2
	National Meteorological Center	2
	National Severe Storms Forecast Center	2
	National Hurricane Center	3
	Reorganization of the National Centers	3
	Office of the Director	3
	Aviation Weather Center	3
	Tropical Prediction Center	4
	Storm Prediction Center	4
	Marine Prediction Center	4
	Space Environment Center	4
	Hydrometeorological Prediction Center	4
	Climate Prediction Center	5
	Environmental Modeling Center	5
	NCEP Central Operations	5
	Relocation of the NOAA Science Center	5
	Numerical Modeling	7
	Regional Systems	9
	Global Data Assimilation and Forecast Systems	11
	Coastal Ocean and Ocean Wave Models	13
	Reanalysis and Climate Data Assimilation System (CDAS)	13
	Coupled Ocean-Atmosphere Model (COAM)	14
	Manual Guidance	15
2.	PLANNED PROCESSING AND TELECOMMUNICATIONS ARCHITECTURE	16
	Standardization	16
	Open Systems	17
	Network Communications	17
	Operating Systems	17
	Programming Languages	18
	Graphical User Interface	18
	Data Formats	19
	High Performance Computing	20
	NCEP Supercomputing Design Objectives and Strategy	21
	NCEP Supercomputing Schedule	23
	Computing Architecture	23
	Hardware	25
	Software	27
	Traditional Mainframe Systems	29
	Disposition of the HDS Systems	29
	Historical Perspective	30

March 13, 1995

Declining role of the HDS systems	31
Approach to Migration	32
Objectives of the Migration from MVS to an Open Architecture	33
Networking and Communications	35
Internal Communications -- Wide Area Networks . . .	35
Internal Communications -- Local Area Networks . .	35
External Communications	36
Data Flow and Data Handling	37
Scientific Workstations	38
Production Control in a Distributed Environment . . .	39
AWIPS	40
Contract Support	41
External Coordination	42
3. RELATIONSHIP TO THE NOAA STRATEGIC PLAN	44
Advance Short-Term Warning and Forecast Services . . .	44
Implement Season to Interannual Climate Forecasts . . .	45
Predict and Assess Decadal-to-Centennial Change . . .	45
4. ALTERNATIVES CONSIDERED	46
5. SECURITY	47
6. ACCESSIBILITY	48
Acronyms and Abbreviations	49

Attachment 1

Statement of Intent between the National Oceanic and
Atmospheric Administration and the National
Aeronautics and Space Administration Regarding the
Development of a NOAA Operations and Research
Center at the NASA Goddard Space Flight Center

Attachment 2

Advances in Operational Numerical Weather Prediction:
An Outlook for the Early Twenty-First Century

Attachment 3

An Overview of Operational Prediction Capabilities at
Fleet Numerical Meteorology and Oceanography
Center

Attachment 4

Memorandum Regarding Massively Parallel Processing at
NCEP

March 13, 1995

Attachment 5

NCEP Wide Area Network Implementation Plan

Attachment 6

Security Plan for NOAA Central Computing Facility
General Support Systems

Attachment 7

Security Plan for the NMC Workstation Facility

Attachment 8

NCEP UNIX Conversion Plan -- DRAFT

1. INFORMATION REQUIREMENTS

The information requirements of the National Centers for Environmental Prediction (NCEP) are extensive, ranging from support for sophisticated research on numerical weather prediction modeling problems to routine word processing necessary for the dissemination of weather forecasts.

In the area of numerical modeling, the NCEP performs research relevant to numerical techniques and model parameterizations designed to develop the world's most advanced operational numerical weather prediction (NWP) models. Research related to these activities explores the application of state-of-the-art technology relevant to supercomputing and massively parallel processing (MPP) systems. It involves extensive experimentation and support for collaborative research projects. If experimentation demonstrates the validity of new techniques, then models that incorporate those techniques are implemented and run routinely, operationally.

In addition to improving the models themselves, NCEP scientists continuously strive to extract the maximum information possible from the output of numerical models. This work includes the development of various techniques for postprocessing model output as well as "ensemble" predictions. The ensemble approach involves multiple runs of a given model from a set of reasonably perturbed initial conditions for the purpose of sampling the range and likelihood of alternative outcomes. Recent experience has shown that these outcomes, when analyzed together, yield forecasts with enhanced predictive ability. Although ensemble forecasting utilizes less highly resolved and/or physically sophisticated models (when compared with current state-of-the-art models) the technique requires large-scale computing systems because of the large number of model runs needed to support it.

Very high-end computing systems are also used to support collaborative research with university scientists and other external researchers. Joint research projects broaden the approach and add new perspective to NWP problems that are of critical importance to modelers and forecasters at the National Centers.

A wide range of pre- and post-processing jobs support the modeling activities of the National Centers. These jobs ingest and format data, store model output, execute secondary model processes, generate derived products, and so forth. These operations have traditionally been performed on mainframe systems.

An extensive amount of data is needed to support the broad scope of operational modeling, routine forecasting, research and development, and individual case studies. The central disk storage capacity necessary for these purposes is on the order of hundreds of gigabytes. Even more extensive data requirements exist for longer range storage. Tape silos and other tape subsystems provide access to off-line storage with a total capacity measured in terabytes.

Extensive networks employing very complex topologies support the National Centers. These networks connect equipment at each site via local area networks; and connect the sites via wide area networks. Networks at the National Centers must support very high transfer rates. Furthermore, as the computing systems connected by these networks grow, so do their communications requirements. The National Centers continuously explore advances in communications technology in order to accommodate their diverse and growing needs. Current networks sustain transfer rates of megabits per second but projected communications requirements are in the range of gigabits per second. In addition, the networks at the National Centers must interface with other networks such as those supporting AFOS, AWIPS, and the OSO Gateway system.

In addition to supercomputers and mainframe systems, the networks at the National Centers support many UNIX servers as well as hundreds of scientific workstations and personal computers. Taken together, these systems comprise a very complex distributed processing environment. Utilizing such network facilities as the Network File System (NFS) and Network Information Services (NIS), this distributed system supports local application development, client-server computing, and scientific visualization methodologies. Researchers, forecasters, technicians, administrators, and managers all use these systems to access computing resources, data, and software as needed, wherever they are found on the network.

Current Organization of the National Centers

The National Meteorological Center presently administers components at three sites. These are the National Meteorological Center (NMC) in Camp Springs, Maryland; the National Hurricane Center (NHC) in Miami, Florida; and the National Severe Storms Forecast Center (NSSFC) in Kansas City, Missouri.

The **National Meteorological Center** develops, produces, and processes meteorological, short-term climate, and oceanographic

March 13, 1995

guidance products and enhanced products. The users of NMC products are National Weather Service (NWS) field offices, the military, and other government and non-government offices. The administrative staff in the Office of the Director handles the personnel, budget, and other administrative and executive matters for the entire NMC.

The **National Severe Storms Forecast Center** maintains a continuous watch for thunderstorm activity in the conterminous 48 states. The NSSFC prepares and distributes outlooks for thunderstorms and forecasts (watches) for tornadoes and severe thunderstorms to NWS field offices, the general public, and other interests. The NSSFC also issues in-flight advisories and hourly bulletins to the aviation industry for hazardous weather phenomena of a convective nature.

The **National Hurricane Center** maintains a continuous watch for tropical cyclones over the Atlantic Ocean, Eastern Pacific Ocean, Caribbean Sea, and Gulf of Mexico from June 1 through November 30. The NHC prepares and distributes hurricane watches and warnings for the general public; and prepares and distributes marine and military advisories for other users. It coordinates with Hurricane Warning Offices, Weather Service Forecast Offices, and Weather Service Offices when a tropical cyclone threatens the United States.

Reorganization of the National Centers

In conjunction with the Modernization and Associated Restructuring (MAR) program within the National Weather Service, the NMC is undergoing an extensive reorganization. In 1995 it will become the National Centers for Environmental Prediction.

The NCEP will evolve over a period of years into an organization comprised of seven service centers, a modeling center, and a group responsible for central operations. The National Centers will occupy facilities in six geographic locations. The NCEP components are as follows.

The **Office of the Director** will be the top-level management of the NCEP. It will be responsible for the formulation and implementation of strategies and policies for all of the components of the NCEP. The Office of the Director will interact with NWS headquarters and with NOAA on behalf of all the National Centers. The Office of the Director will coordinate research, system operations, and analysis and forecasting throughout the National

March 13, 1995

Centers to ensure the highest quality of products and services. It will perform broad administrative duties for the Centers and, except in the case of the Space Environment Center, will be the final authority regarding personnel and budgetary recommendations. The Director of the Environmental Research Laboratories (ERL) within the Office of Oceanic and Atmospheric Research (OAR) will be the final authority on these matters for the Space Environment Center. The Office of the Director will be located at the NOAA Science Center in Camp Springs, Maryland.

The **Aviation Weather Center** (AWC) will be located in Kansas City, Missouri. It will issue operational aviation analyses and forecasts within domestic and international airspace and will have the primary responsibility for global aviation products valid up to 48 hours into the future. The AWC will be organized from components within the National Centers that are currently located at the NSSFC in Missouri and the NMC in Maryland. The AWC will be housed in a new facility planned for construction in Kansas City.

The **Tropical Prediction Center** (TPC) and the National Hurricane Center will continue to be located in facilities in the Miami, Florida area. This Center will issue a range of products that includes cyclone watches and warnings as well as tropical weather guidance and forecasts. These products will cover the tropical regions of the Western Hemisphere from the equator to 30 degrees north latitude and extend from 0 to 5 days. A new facility that will house TPC/NHC and the Miami Weather Forecast Office (WFO) is under construction on the campus of Florida International University.

The **Storm Prediction Center** (SPC) will be located in Norman, Oklahoma and will consist largely of NMC components presently located within the NSSFC in Kansas City. The SPC will provide guidance forecasts on hazardous weather, guidance outlooks, and watches. These products will cover the conterminous United States from 0 to 24 hours. The SPC will occupy part of a new facility that will be constructed on the campus of the University of Oklahoma. This building will also house organizations such as the National Severe Storms Laboratory and the Norman WFO.

The **Marine Prediction Center** (MPC) will be created in Monterey, California. This Center will issue analyses, guidance, forecasts, and warnings for the marine boundary layer and the ocean surface. These products will cover the high seas as well as the offshore and coastal areas of the Atlantic and Pacific Oceans; and will extend from 160 degrees east to 35 degrees west longitude, north of 30 degrees north latitude. The time scale will

March 13, 1995

range from 0 to 5 days. The MPC will include both NCEP and National Ocean Service (NOS) personnel. The NCEP portion will be formed from components of the Meteorological Operations Division (MOD) of the NMC that will relocate to Monterey. A new building will be constructed in California for the MPC.

The **Space Environment Center** (SEC) is located in Boulder, Colorado. The SEC is and will remain a part of NOAA/OAR/ERL and is also a component of the National Centers. Its mission is to serve the nation's need to reduce the adverse effects of solar-terrestrial disturbances on human activities. It provides global forecasts and summaries of solar-terrestrial conditions out to a period of several weeks.

The **Hydrometeorological Prediction Center** (HPC) will be located at the NOAA Science Center in Maryland. It will provide analyses, guidance, and ready-to-use forecasts for a range of hydrometeorological parameters over the United States out to 7 days. The HPC will consist of the parts of MOD that are not moved to the AWC and MPC; and will include the National Precipitation Prediction Unit formed jointly by the National Weather Service and the National Environmental Satellite Data and Information Service (NESDIS).

The **Climate Prediction Center** (CPC) will also be located at the NOAA Science Center and is the continuation of the Climate Analysis Center of the NMC. It will provide climate analyses, forecasts, and monitoring products. This wide range of products will variously cover North America, the Northern Hemisphere, and the entire globe. The time frame for CPC products will extend from week-two (8-14 days) out to multiseason forecasts.

The **Environmental Modeling Center** (EMC) will be responsible for the development of numerical models to perform analyses and forecasts of the atmosphere and oceans. The spatial scales range from mesoscale to global; the temporal scales range from hours to multiseason. The EMC is the successor to the Development Division of the NMC and will be located at the NOAA Science Center.

NCEP Central Operations (NCO) will operate the large-scale computing facilities of the NCEP; and will run the numerical models that are central to the functioning of the National Centers. NCO will also provide the communications, software, and support infrastructure that connects and integrates the other Centers. The NCO will evolve from the Automation Division of the NMC and will have components at two sites in Maryland -- the NOAA Central Computer Facility (NCCF) in Suitland and the NOAA Science Center.

March 13, 1995

Relocation of the NOAA Science Center

The NOAA Science Center refers to the aggregate of the NOAA components presently located in the World Weather Building in Camp Springs, Maryland. For a number of reasons, these components are proposed to relocate to the NASA Goddard Space Flight Center (GSFC) in Greenbelt, Maryland and to be renamed the NOAA Operations and Research Center. A copy of the draft Statement of Intent between NOAA and NASA is Attachment 1 to this document.

As discussed in the Statement of Intent, the proposed relocation presents a unique opportunity to assemble appropriate talents, skills, and disciplines from NOAA and NASA in a way that will make considerable contributions to our country's effort to define and maintain the planet's environment. It is in the best interest of the U.S. Government to foster U.S. world leadership in global change research, associated research and applications efforts, and to make the best use of available resources and opportunities. The assemblage of a joint NASA/NOAA capability would support these interests. For example, this would facilitate joint NASA/NOAA efforts in the U.S. Weather Research Program recently inaugurated by the Federal Coordinating Council on Science, Engineering, and Technology. By promoting maximum scientific exchange among NASA and NOAA scientists, the U.S. public would directly benefit from an accelerated improvement of NOAA's operational prediction and assessment mission and from the joint research related to NASA's Mission to Planet Earth program.

Such a collocation would also result in a number of economies accruing to NOAA by being able to draw on NASA's central plant resources. The basic infrastructure is already in place so only a tie-in to existing systems is needed. The available utilities that could be extended to this new facility include: chilled water, steam, potable water, sewer, power, lighting, and communications. Back-up chilled water and power can be provided by expanding the existing central plant systems.

The relocation to GSFC will include the supercomputer and main-frame systems presently at the NCCF and so the economies noted above will be real and substantial. However, it is the synergistic effect of bringing together NOAA and NASA environmental scientists that offers the greater benefit to the nation. The aim of collocating these groups is consistent both with the modernization program of the National Weather Service and with the NOAA Strategic Plan.

March 13, 1995

All permits for construction and operation have been obtained including wetland mitigation, archeological survey, water management, rights of way, and emission.

Also, the site is fenced-in with guardhouses at the points of entrance. Nearby dining facilities for 1500 persons in the NASA Earth System Science Building are also planned with the possibility of satellite food areas. A conference center for up to 250 persons including media and communication facilities is planned as well. Child care, a credit union, and a fitness center/health unit would also be available for use.

It is proposed that the above topics be discussed after the Statement of Intent is signed and a joint NASA/NOAA Steering Committee for the Development of the NOAA Operations and Research Center at the GSFC is established. It is important to emphasize that what is proposed in the Statement of Intent is *collocation* of the NOAA Operations and Research Center at GSFC, not a *consolidation*. NOAA's mission is distinctly different from NASA's which precludes the possibility, for example, of NASA providing computer back-up for NOAA's real-time, operational forecast system. On the other hand, such a collocation would significantly enhance the scientific interaction of NOAA and NASA personnel, help to facilitate technology transfer, help to ensure that satellite data obtained from NASA's Mission to Planet Earth are available to the Earth science community over the long term, and sustain the preeminence of the U.S. in the use of satellites and other advanced technologies for studies of the Earth and its environment.

The NOAA components presently under consideration for collocation are:

- | | | |
|--------|---|---|
| NWS | - | National Centers for Environmental Prediction
(Washington D.C. based components) |
| | - | Techniques Development Laboratory, Computer
Systems Section |
| | - | Hydrometeorological Research Laboratory |
| | | |
| NESDIS | - | Office of Research and Applications |
| | - | Office of Satellite Operations Control Center |
| | - | Office of Satellite Data Processing and
Distribution |
| | - | Satellite Data Services Division |
| | | |
| OAR | - | Air Resource Laboratory |

March 13, 1995

Numerical Modeling

Progress in the science of numerical weather prediction (NWP) has occurred as a result of three factors: advances in our ability to observe the atmosphere, advances in our ability to understand and model the behavior of the atmosphere, and advances in our ability to compute the future evolution of the atmosphere, based on our observations and models (McPherson, Attachment 2.)

The numerical modeling plans for the National Centers are designed to take advantage of the union of the three factors cited by Dr. McPherson in the above quote. Those plans assume that the computers used in NWP will continue to offer increased performance -- an increase from about 1 billion floating point operations per second in 1990 to 1 trillion floating point operations per second early in the next century. Attachment 2 also provides an overview that explains why the National Centers are pursuing a computing strategy that involves higher resolution models in the mesoscale; and ensemble (statistical) techniques for longer-range forecasts. This typifies an approach that matches the computing requirements associated with different forecast problems to different classes or configurations of computing systems.




Figure 1

This figure demonstrates 40 years of progress in the use of operational numerical weather prediction models at the NMC. It charts a specific measure of forecast skill over time and also identifies some key events to the time line.




Figure 2

This figure represents the reduction over the last 15 years of a specific measure of the error in NMC numerical model forecasts of upper level winds.

The discussion that follows provides an overview of numerical weather prediction modeling divided into several classes or forecast systems. This is a comprehensive but not exhaustive list of such systems. All references to dates are approximate and reflect current estimates as to the availability of advanced computing systems and the results of diverse research and development activities.

Regional Systems

Nested Grid Model (NGM)

The NGM is a 16-layer model with 80 km resolution. It generates 48-hour forecasts twice a day and is used for model output statistics guidance. Development has been frozen on the NGM since 1990 and the model itself will be discontinued in 1998. This is after the Class VIII system and the model enhancements that it will support are in operational use.

Early ETA Forecast Model (ETA)

The early ETA run involves a 38-level model with 80 km resolution. Like the NGM it generates 48-hour forecasts twice per day. The ETA utilizes a first guess from the global data assimilation system (GDAS) and an ETA optimal interpolation (OI). The need for the early ETA run will be obviated by the aviation (AVN) global model when it begins to run four times per day in 1996. Once the AVN demonstrates accuracy comparable to or better than the ETA for precipitation guidance in the 24-48 hour range, then the early ETA will be discontinued.

Mesoscale ETA Data Assimilation System (EDAS)

This OI-based data assimilation system is used to initialize the mesoscale ETA model. It is performed on the model domain at 29 km resolution and for 50 levels. It will become operational in early 1995 and will initially run two 12-hour cycles per day, each with an intermittent (3-hour) assimilation of data. Later in 1995 the EDAS will begin to run four times per day. In 1998 the EDAS will evolve into a four-dimensional variational regional data assimilation scheme based on the adjoint of the ETA model. It will then run four times per day with continuous data assimilation over 6 hours for each of the four cycles.

Mesoscale ETA Forecast Model (Meso-ETA)

This is a 50-layer model with 29 km resolution. Initial conditions come from the EDAS. The Meso-ETA will become operational in early 1995 and will produce 33-hour forecasts twice per day. The base times for these runs will be 03Z and 15Z in order to use updated boundary conditions from the AVN model. Later in 1995 the Meso-ETA will begin to run four times per day. Special ETA forecasts will be run over Alaska if computationally feasible. After the arrival of the Class VIII system, plans are to consider an increase in the horizontal resolution

of the Meso-ETA to 15 km and to increase the number of levels to 70.

Regional Spectral Model (RSM)

This model is presently under development. The RSM, now run experimentally within the global spectral model, is a 28-level model with 40 km resolution covering a North American sub-domain of the NGM. It produces one 48-hour forecast per day for comparison with the ETA model. The RSM should become operational in 1995, perhaps with an increase to 42 levels. After the operational implementation of the Class VIII system in 1998 the RSM will run with 20 km resolution within the AVN model four times per day as an enhancement over North America. In the same time frame a non-hydrostatic version with 10 km resolution will come into use experimentally.

Hurricane Model

The Quasi-Lagrangian Model (QLM) is currently operational. This is an 18-layer model with 40 km resolution. It produces 3-day forecasts when needed. The GFDL Multipynested Movable Mesh (GMMM) model, also an 18-level model but with 20 km resolution on an inner grid, is now running experimentally. The QLM will probably be discontinued in favor of the GMMM in 1995. The GMMM will increase its resolution to 10 km when it moves to the CLASS VIII system in 1998.

Rapid Update Cycle (RUC)

The RUC is a hybrid sigma-isentropic analysis and forecast system. It utilizes a 3-hour, OI-based data assimilation cycle to produce 12-hour forecasts eight times per day. This is a 25-layer model with 60 km resolution. With the implementation of the Class VIII system in 1998 the RUC will run hourly, 24 times per day. The resolution of the model will increase to 50 levels and 30 km in that same time frame.

Regional Systems Outlook

The outlook for the 5 years after 1998 is that the NCEP will run a national domain mesoscale model at 5-10 km resolution based on a non-hydrostatic version of the ETA model (possibly the RSM), together with ensembles of lower resolution. This system will be coupled with multiple storm-scale (1 km or less) models used for very short range forecasts (up to 6 hours). The storm-scale forecasts will be nested within the national domain in areas threatened by severe weather. Together with extensions of NCEP models, candidates for such storm-scale

March 13, 1995

models could be those developed at universities or other research centers. Candidate models will be evaluated for potential operational use in the Model Test Facility to be established within the Environmental Modeling Center.

Global Data Assimilation and Forecast Systems

Aviation (AVN) Global Model

The AVN is a 28-level model with 100 km horizontal resolution which runs twice per day producing 72-hour forecasts. It will be refined to 40 levels and approximately 80 km resolution in the summer of 1995. Another advancement will be the conversion from spectral to semi-lagrangian methodology in late 1995. The AVN will run four times per day beginning in early 1996. In 1998 the AVN will run with 60 km or finer resolution using 50 levels. The RSM will run within this model at 20 km resolution providing the main synoptic guidance over North America.

Medium Range Forecast (MRF) Model

The MRF is the same numerical model as the AVN but it is run under a wider variety of conditions and for different purposes. In late 1994 it began to run once per day at the resolution of the AVN to produce a 7-day forecast. A low-resolution version of the MRF (using a 2.5 degree grid) runs out to 16 days as a tool for ensemble forecasting. The low-resolution MRF will also run experimentally on the Class VII machine with an imbedded version of the RSM as referred to above.

In 1998 the MRF will probably be used in ensembles of a few runs with 70-100 km resolution out to 7 days; and ensembles of many low-resolution runs for days 7 to 30.

Ensemble Forecasting

The encouraging experience thus far at the NCEP with the ensemble system presently available for the 1-5 day global forecasts suggests that this system can be improved and extended to longer ranges. The use of ensembles of forecasts has already begun and will provide a quantitative foundation for probabilistic synoptic forecasting beyond the first week. Even for short-range forecasting the use of ensembles will provide guidance for the probability distribution of precipitation and extreme events.

Global Data Assimilation System (GDAS)

March 13, 1995

The GDAS will continue to use the spectral statistical interpolation technique but will add refinements as they are developed. By 1998 it will apply four-dimensional variational methods with a 12-hour assimilation interval and continuous data utilization.

The GDAS will also use a number of new or expanded data sources. Satellite instrumentation in particular will provide new data that will be incorporated by the GDAS over the next several years. Data from profilers and from new aircraft sensors will be assimilated too. In addition, the GDAS will be refined to better utilize information from traditional sources such as pressure tendencies and cloud reports from surface observations.

Global Systems Outlook

The outlook for NCEP global models after 1998 is that they will have a horizontal resolution of about 20 to 40 km and will run through 72 hours in the case of the AVN and through 7 days for the MRF. Ensembles of somewhat lower resolution model forecasts will be run for one month or longer. With better data assimilation, skillful tropical forecasting based on these models and ensembles may become possible through 10 days. Ozone forecasts should be routinely available and forecasts of atmospheric contaminants may also prove useful.

Coastal Ocean and Ocean Wave Models

Ocean Wave Model

The third generation WAM (**W**Ave **M**odel) global model began to run on the Class VII system in late 1994 replacing the NOAA Ocean Wave (NOW) model. It runs twice a day out to 72 hours with 2.5 degree resolution. By 1998 NCEP will utilize a nested grid ocean wave model on the Class VIII system. The resolution of this model will be one degree in the open ocean and 0.25 to 0.5 degree in selected regions along the East and West Coasts, the Gulfs of Alaska and Mexico, and in the area around Hawaii.

Coastal Ocean Data Assimilation and Prediction System (CODAPS)

Preliminary work has been performed relative to the development of a prototype Mellor model and data assimilation system. This is designed to run routinely for the East Coast of the United States. If this experiment is

successful, it will become operational in 1996 or 1997 and a similar experiment will begin for the West Coast.

Ocean Systems Outlook

In the out years of 1999 and beyond, high resolution (5 km) coastal nowcasting and short-range forecasting should provide guidance to fisheries, ships, and environmental monitors along the entire North American coast with forecasts of thermal structure, currents, and water levels, including tides and wind induced storm surges.

The wave models over the coastal domains will be coupled to the CODAPS models to take into account wave-current interactions which significantly modify the waves.

Reanalysis and Climate Data Assimilation System (CDAS)

Reanalysis

The reanalysis system will use 28 levels and 2.5 degree resolution. It will reanalyze 35 years of data, performing at a rate of 30 days of analyses per calendar day. Operational execution of this program began in late 1994. After the implementation of the Class VIII system in 1998, another reanalysis cycle will begin with a domain of approximately 40 levels and one degree resolution.

CDAS

This program is designed to monitor the climate and climate change. The CDAS will become operational coincident with the reanalysis system. It will assimilate current data at the same resolution as the reanalysis program. The CDAS and the reanalysis system will both run on the same computer system.

Reanalysis and CDAS Outlook

In the period after 1998 a long reanalysis with a state-of-the-art (but frozen) data assimilation system will be performed about every 5 years. This will provide guidance to climate monitoring and forecasting efforts. It will also be valuable to compute "perfect prog" statistics, linking model analysis with station observations using an extremely long training period. This can be combined with adaptive model output statistics in which the model forecast is used to predict the analysis, requiring short training periods.

Coupled Ocean-Atmosphere Model (COAM)

The COAM uses an 18-level version of the MRF model with approximately five degree resolution. This is coupled with a Pacific Ocean model to produce long-range forecasts. The Coupled Model Project will use the Class VII system to make one 12-month coupled model forecast per week. This forecast permits dynamic interaction between the atmosphere and the ocean, a process that currently results in significant systematic biases in the forecast of sea surface temperatures. These biases are corrected in a post-processing step. The Coupled Model Project will also produce two 6-month forecasts each week from the atmospheric model alone using the corrected sea surface temperatures from the coupled model run. These forecasts are both for the same 6-month period and are run from different initial states.

By late 1995 the COAM will be extended to a 28-level 2.5 degree resolution MRF/Pacific Ocean model. The methodology will be essentially unchanged from the above and will produce one 12-month coupled model forecast and two 9-month atmospheric model runs, again using corrected sea surface temperature forecasts, each week.

When the COAM moves to the Class VIII system in 1998 it will be enhanced by the inclusion of a global ocean model. It will also couple with the regional spectral model to provide higher resolution over the United States. By this time it will produce one fully coupled 12-month forecast per week. It is anticipated that in this time frame the model biases in sea surface temperature forecasts will have been largely removed and that the dynamic processes of the coupled model will yield superior results.

In the period after 1998, routine 1-year coupled model predictions will be produced which are expected to have significant skill in global variations associated with the El Nino/Southern Oscillation (ENSO) phenomena. Initial conditions for such forecasts will come from a global ocean/atmosphere time-continuous data assimilation system.

Manual Guidance

March 13, 1995

In addition to developing and executing numerical models and managing the computing and communications systems upon which they rely, the National Centers provide a wide range of manual forecast and guidance products. These products include text, graphics, and gridded numerical data and extend from very short-term severe weather forecasts to multiseason climate outlooks. The preparation of these manual guidance products is inherently very time consuming and labor intensive. Their preparation is, in fact, the chief activity within each of the National Centers excepting only the NCO and EMC. The creation of these products is essential so that the National Weather Service can present a comprehensive and coordinated set of forecasts to national and international users. It is the intent of the NCEP to use modern computing and communications technology, an open systems approach, and a small number of well-defined standards, all consistent with the AWIPS program, to create a comprehensive program for the creation of manual forecast and guidance products

2. PLANNED PROCESSING AND TELECOMMUNICATIONS ARCHITECTURE

The overall processing and telecommunications architecture that supports the National Centers today is a very complex one. It encompasses supercomputers, traditional mainframe systems, scientific workstations, wide area networks (WANs), local area networks (LANs), multiple operating systems, and several communications protocols. The complex information requirements of the National Centers necessitate a complex architecture. However, this architecture can and should be designed so that its complexity is manageable, controllable, consistent, reasonable, comprehensive and understandable. This goal can be accomplished by following two principles.

The first of these principles is this: to the greatest extent possible the architecture should be planned and designed to anticipate general requirements. It should not evolve as a conglomeration of ad hoc responses to immediate and specific needs. Neither should the system architecture indefinitely accommodate systems that were developed in the past to solve specific problems. Such systems often have a limited focus, function within a narrow domain, and cannot interface easily with other systems.

Second, the architecture should evolve with a high regard for carefully considered, unambiguous standards. These standards themselves will certainly evolve, but under the proper circumstances the overall architecture for processing and communications at the National Centers will become a system comprised of mutually supportive subsystems. The result will be a robust, comprehensive, understandable, and manageable system.

The appropriate circumstances are these:

- The NCEP must identify areas in which standards can be defined.
- The National Centers must strive to select appropriate standards.
- Managers within the NCEP must consistently take advantage of and enforce those standards.

The National Centers are firmly committed to these principles.

Standardization

March 13, 1995

Over the last several years a number of relevant standards have emerged at the National Centers. For the most part these were not the result of any specific, conscious process. Rather, their adoption was due largely to the fact that scarce resources have forced a narrowing focus as to what is supportable. In any case, it is now possible to posit a small number of such standards that can clarify the nature of the overall planned processing and telecommunications architecture of the National Centers. These standards fall naturally into just a few general areas.

Open Systems

The overarching purpose of the standards that follow is that of establishing and maintaining an open computing environment at the National Centers. The complexity and diversity of the NCEP requirements is such that they demand a wide variety of capabilities and equipment. In order for such an environment to be manageable, it must be possible to move many, even most, applications about the network, from site to site or machine to machine, to take advantage of appropriate resources. This is only possible if the hardware and software that comprise those resources share fundamental characteristics. These characteristics are defined by the standards adopted by the National Centers and as a group they establish an open computing architecture.

Network Communications

The predominant means of network communications within the National Centers today is Transmission Control Protocol/Internet Protocol (TCP/IP) ethernet. Many other network architectures and communications protocols are in use at the National Centers and these cannot all be discontinued immediately. However, these alternatives are either proprietary or are based on older technologies and accordingly their use is in decline. When possible, as legacy systems are upgraded, the upgrades will utilize TCP/IP ethernet.

Two standards consistent with TCP/IP ethernet offer specific additional network services to the National Centers -- NFS and NIS. NFS is an independent network service that supports remote access to disk subsystems. NIS is a distributed name service that simplifies the administration of a network of computer systems. NIS maintains a database of user names and translates those names to network addresses. Both NFS and NIS are operating system and machine-type independent.

As new communications standards emerge, the National Centers will monitor their applicability. One emerging standard that is of particular interest is Asynchronous Transfer Mode (ATM). As ATM standards are defined and as vendors bring ATM products to market, the National Centers plan to develop prototype ATM network

capabilities to test the feasibility and practicality of its use. Because ATM can encapsulate TCP/IP ethernet, there is a natural growth path between the two.

Operating Systems

Although the UNIX operating system comes in many varieties, there is a basic compatibility and extensive commonality between any two UNIX implementations. Those characteristics transfer to all of the UNIX systems on a network. This inherent compatibility has resulted in a general migration to UNIX throughout the National Centers. Except for most office automation tasks, UNIX has become the standard NCEP computer operating system. Legacy systems that run other operating systems, such as MVS, will be phased out and the replacement systems will run UNIX.

Programming Languages

FORTRAN has long been the standard programming language employed at the National Centers. For scientific applications this will continue to be the case. FORTRAN is itself evolving and the National Centers will strive to maintain code in the current ANSI standard for the language. FORTRAN77 is in current use but FORTRAN90 is beginning to supplant it.

Another language is commonly used in the UNIX programming environment and that language is C. C is particularly appropriate to computer graphics and networking applications in UNIX, which is itself written in C. As the National Centers have increasingly distributed computing over various UNIX systems, C has gradually become a second standard programming language among scientists and programmers. Software development at the National Centers will rely upon applications written in standard, portable versions of C and FORTRAN.

In most cases vendor specific language extensions will not be employed. However, the application of vector processing systems, multiple central processor unit systems, and MPP systems to large NWP models are, for now, exceptions to this principle. Running large models on high-end systems requires that codes be optimized to specific system characteristics. In some cases this severely limits their portability. However, the nature of these large special-purpose programs is such that they need the unique capabilities offered by the high-end systems in order to be of practical use. Therefore they are inherently not portable to other network devices anyway.

Graphical User Interface

A key component of a distributed computing environment is the graphical user interface (GUI). A comprehensive, well-developed GUI shields the user of such a network system from the complex

processes that support his or her application of the system to the performance of specific tasks. In the perception of many users, the GUI is the system. Quite a number of GUI products are available commercially but most have severe limitations. They are either dependent upon proprietary software libraries or have limited applicability -- some are specific to particular hardware platforms, for instance. There are, however, two basic tool kits that comprise the standards upon which a UNIX GUI can be built. These are X Windows and Motif.

The X Window System is based on a client-server model and provides a standard windowing environment for applications programmers. It was developed at the Massachusetts Institute of Technology and is supported on all UNIX platforms. Motif is a GUI standard designed for X Windows by the Open Software Foundation. It consists of a widget set, a window manager, a style guide, and a user interface description language. Together, X and Motif comprise a comprehensive, portable, vendor-independent set of tools for graphical presentations and they are the relevant standards used by the National Centers for GUI development.

Data Formats

Consistent with the establishment of computing and communications standards, the National Centers have adopted two standard data formats. Both are approved by the World Meteorological Organization (WMO). The first of these is GRIB (**GR**idded **B**inary), a general purpose, bit-oriented data exchange format. It is an efficient vehicle for transmitting large volumes of gridded data between meteorological centers over high-speed telecommunications lines using modern protocols. The second is BUFR (**B**inary **U**niversal **F**orm for the **R**epresentation of meteorological data). BUFR is a binary code designed to represent any meteorological data by employing a continuous binary stream. The code form may be applied to any numerical or qualitative data type.

Users at the National Centers have come to accept the value and utility of these standard data formats and today most NCEP data is stored and transmitted in them. The National Centers will use GRIB and BUFR for data exchanges between Centers. These standard data formats will have a central role in the further evolution of the computing architecture at the NCEP for the foreseeable future.

The standards noted above are not very remarkable for a complex computing environment. Most large organizations have adopted the open systems architecture approach and most of these standards are commonly employed for that purpose. Attachment 3 is a very

March 13, 1995

brief account of a comparable approach taken by the Fleet Numerical Meteorology and Oceanography Center of the United States Navy. The idea of a gradual transition from legacy systems inconsistent with this environment to open systems is a concession to the realities of the cost and effort involved in such a transition.

Given these principles and standards, the fundamental computing architecture at the National Centers is, and will remain, a multi-tiered one. The explicit layers of this architecture are as follows.

- High-end supercomputers or specialized MPP systems are required for advanced numerical models and such models cannot be run on anything less.
- Servers and other special-purpose machines are needed to perform specific pre- and post-processing tasks.
- Scientific workstations enable users to interact with other systems, including supercomputers, to display model output and other information graphically and to request specific services from other network devices as needed.
- All classes of NCEP computers are compatible and are accessible to users via TCP/IP ethernet.

The principal subsystems of the overall processing and telecommunications architecture for the National Centers are discussed under the sub-headings below.

High Performance Computing

NCEP high performance processing is driven by specific programmatic needs and requirements. Paramount among these is the very high computational capacity needed to implement continuously improving NWP techniques into a scheduled operational environment that meets firm product delivery deadlines. A secondary factor is the proliferation from various sources of observational data. Another significant requirement is the demand for new analysis and forecast products from the end users.

As noted earlier, this is a period of rapid progress in the science of weather forecasting. Continued support for contempo-

rary meteorological research and the operational implementation of its conclusions necessitates continued growth in the computer capacity that helps make it possible. Providing this computer capacity enjoins a careful balancing of several aspects of available technology to produce a smoothly integrated and functioning system. A balanced NWP forecasting system must include a flexible processing infrastructure as well as powerful CPUs. The infrastructure of such a system requires hardware and software components that interact with the central compute engines to provide a full repertoire of products and services as discussed below.

The state of development of NWP models at any point in time is such that the demands of individual models vary considerably. Thus the computational requirements of the models can best be met by a range of computer systems. The National Centers have responded to this state by deliberately employing a strategy that causes an overlap at the NCCF with regard to high performance computers. The augmentation of high performance computer systems for operational use has consistently benefited from the retention of the previous generation system while the NCEP strides forward to the next. This has assured users and customers of the operational stability necessary to maintain existing services while providing a platform on which to develop the next generation of operational systems.

Aside from major new acquisitions, the computer systems used by NCEP are in a continuous state of improvement and enhancement. High-end computers characteristically do not deliver their rated, nominal performance. Those responsible for such resources repeatedly adjust system configurations and augment subsystems to attain better performance. As an example, the performance of the Cray C90 system at the NCCF is rated at approximately 16 billion floating point operations per second (gigaflops, or GFLOPS). The actual performance on codes at the NCEP -- codes that employ a broad range of mathematical functions and deal with practical issues like getting data in and out of the machine (conditions often overlooked in theoretical performance claims) is closer to three GFLOPS on current operational models but as high as seven GFLOPS on some experimental models. Throughout the life cycle of this system, the NCEP and the vendor will strive to adjust and tune hardware and software subsystems in order to deliver more performance in real-world applications.

NCEP Supercomputing Design Objectives and Strategy

Mid-range Processing

March 13, 1995

NCEP agreed to an engineering change to the Cray C90 contract in September, 1994 that exchanges the Cray Y-MP/864 and the Cray Y-MP/EL2 systems for a pair of Cray J916 computers and supporting disk sub-systems. This transaction enhances the support role provided by these systems while also providing a migration platform for some of the NCEP's general mid-range computational needs. This mid-range processing includes some work presently performed on the HDS MVS systems. It specifically does not include the majority of NCEP's currently operational numerical models. However, some older models can run in this environment and it will be used for several specific purposes. Three examples of this are the reanalysis project, the provision of backup capability in the event of degraded operations, and limited research projects. Mid-range processing will be addressed further during the next supercomputer procurement or in a separate procurement action.

System Life Cycle

System upgrade cycles are determined by improvements in NWP modeling as constrained by what is available in the marketplace and by funding. There is no predefined workload, such as might be encountered in traditional business processing, that drives an upgrade cycle. The objective of the NCEP in this regard is to apply the best available computing architecture to meet scientific needs when reliable computational speed can improve the quality of the results.

Replacement Timing

Industry technology and commercial offerings are central motivating factors in the selection of these high-end systems. The availability of proven technology, with reasonable and acceptable cost, helps to determine the replacement cycle. The National Centers attempt to upgrade at intervals that provide not less than five times the extant level of performance. NOAA has traditionally been near the forefront of supercomputing, relying on the latest available high-end systems that use proven, low-risk technology. There has been consistent and enormous progress in high performance computing since the early 1960's. In that time, classes of numerically intense computer systems have repeatedly become almost commodity products, available as off-the-shelf items. The maturation of this industry is illustrated and reflected by its foreshortened system life cycle period. The pace of progress has accelerated such that system replacement cycles of the type characterized above are less than five years. Because the NCEP uses an alternating operational platform strategy to minimize the implementation risks for operational codes, this results in nearly continuous acquisitions being undertaken by the NCEP.

March 13, 1995

Replacement Strategy

NCEP acquisitions are based on available performance as moderated by the need to minimize risk. The industry in the mid-1990's is at a watershed with regard to supercomputer architectures. A number of competing architectures are being tested in the marketplace. Each design offers merit for a particular computational environment. It may be that certain configurations will prove inappropriate for the kinds of models employed in NWP. However, several designs show promise and they will be evaluated by the NCEP. The industry appears to be on the verge of a breakthrough regarding MPP systems that could be quickly followed by the commercial availability of robust, low-risk massively parallel systems. The NCEP will monitor developments in this area and assess the associated risks.

Industry Partnership

For many years the National Centers have made the source code for its operational forecast models and other representative benchmark programs available to the computer marketplace. This allows computer designers to make use of these codes to develop systems using complex real-life applications. The NCEP benefits directly from this work and willingly involves a few scientists in it. They modify and restructure working models to run on new, evolving systems. Government scientists optimize programming constructs to test systems under current manufacture and offered in the marketplace. The NCEP has established this partnership with industry to assure that computer designs are suitable for its needs and applications. In the course of this work, NCEP gets to test processes on a wide variety of systems and to consider new paradigms and new programming techniques. In this way, the NCEP ensures that acquisition specifications are open to many architectures while still addressing the unique nature of numerical weather modelling.

NCEP Supercomputing Schedule

Acquisition schedules are predicated upon system life cycles as driven by the availability of cost effective commercial, off the shelf (COTS) hardware and software. In particular, the high performance computer systems at the heart of NCEP's operations are acquired only after suitable platforms evolve sufficiently to support stable operations. The NCEP has been able to periodically replace systems when their residual value drops below their maintenance costs. Normal progress within the industry presently results in new comparable hardware at significantly reduced initial and operating costs about every 3 to 5 years. The re-engineered Cray C90 contract is a case in point. The newest generation of air-cooled, power conserving systems have near parity in performance with the Cray Y-MP and offer a significant

March 13, 1995

reduction in both acquisition and operating costs. The J90 systems are about 4 years behind the Y-MP in terms of their availability in the marketplace. Obviously, a somewhat different set of conditions prevail at the pinnacle of performance.

The tentative time frame for supercomputer upgrades at the National Centers is presented in the following table.

Task	Time Frame
NCEP Architecture Plan Completed	Early calendar year 1995
Install Cray J916s + NDA	Mid calendar year 1995
Procure additional near-line storage capacity	Late calendar year 1995
Improve intersystem communications	Early calendar year 1996
Procure and install experimental MPP system	Mid calendar year 1996
Begin preparation of C90 upgrade RFP	Mid calendar year 1996
Issue C90 upgrade RFP	Early calendar year 1997
Install new supercomputer system	Mid calendar year 1998

Computing Architecture

One key to minimizing risk is diversification. In a sense, the entire High Performance Computing industry is diversifying. For example, Cray Research, a longtime high-end computer manufacturer, currently has offerings in both traditional supercomputers and in the emerging field of massively parallel systems. As a consumer of this technology, NCEP recognizes that this is an extraordinary period of change in the high performance computer marketplace. With rapid change, however, comes the risk of choosing the incorrect way to change. There are a number of alternative solutions to the supercomputing needs of the NCEP, a few of which are viable, but some of which will play no part in NCEP's future. The prudent course of action in the face of rapidly changing options is to employ pilot projects for purposes of assessment, using NCEP benchmarks and models. Such pilot projects have been undertaken by the Environmental Modeling Center. The results to date have been extremely promising with regard to parallelizing NWP code and running it on equipment available from several vendors. As an aid to decision making processes, these pilot projects will continue to be supported as an investment in appraising and evaluating this emerging technology.

The next step in this pilot project strategy is to obtain an MPP system for internal use at the National Centers. The purpose behind the acquisition of such a system, planned for 1996, is to use it in conjunction with other resources, not as a replacement for other systems. Attachment 4 is a proposal that describes how an MPP will be used for numerical modeling.

March 13, 1995

Some of the competing high-end technologies in the marketplace today are described next.

Monolithic Parallel Vector Processors

The traditional supercomputer used by the NCCF is constructed as a small collection (typically eight to sixteen) of powerful vector processors that can be used to multiprogram several batch jobs simultaneously. The speed of computation derives from the capacity of each CPU, and system performance is achieved by overlapping task execution among several CPUs. Current operational NCEP models use up to fifteen of the sixteen CPUs in the Cray C90. (The last CPU is reserved for support functions and to ensure timely model completion. In the event that one CPU fails, the spare can be applied to the application). In good implementations of this approach, a compiler can perform fine-grained microtasking decompositions at the "DO LOOP" level, thus improving the degree of parallelism.

Massively Parallel Processors

This emerging technology is the most significant recent development in the field and seems to provide the greatest promise for NCEP in the longer term. Consisting of up to 2048 processors linked together via shared memory or a high speed internal data network, these systems divide computational work into the concurrent execution of interdependent parts. Some applications are more readily decomposed in this fashion than are others. The granularity of task division is typically finer than in other methods consisting of threads of executable elements. As with the microtasking technique of traditional supercomputing, threads inherit a common memory domain as peer elements.

Clusters or Gangs of Low Cost Workstation-level Processors

A popular configuration, used by Silicon Graphics Incorporated and some other manufacturers, this method couples multiple processors into a single frame and under a single operating system. Effective operation depends on inter-node communications within the system. This approach is being carefully watched by NCEP because, although not presently useful for high-end forecast models, it offers some promise for certain kinds of simulation servers and statistical analyses, as well as for pre-processing and post-processing functions of the kind traditionally done on mainframes.

Parallel Vector Machines

Another relatively new approach, this technology distributes the workload in a manner similar to those described above, but at a significantly coarser level. Only discrete task-level elements can be parallelized due to the need to distribute executable elements over a high speed network. Usually, the processing

nodes are physically removed from each other in this approach, but some implementations have them collocated. At the moment, this design is not being pursued by NCEP because of the lower performance range they provide and because of the very high network bandwidth they require to achieve effective results.

Hardware

There are two drivers for the choices that NCEP makes regarding the hardware for high performance computing. These are industry capacity and price/performance comparisons. These considerations drive the development of applications as well as system implementation and deployment.

It is well recognized in the scientific community that weather forecasting is one of the nation's Grand Challenges in the sense that the effort necessary to meet it demands ever more powerful computers. Weather forecasting is one of the most difficult large-scale problems actively pursued by numerical modelers. Adding to the conventional designation regarding numerically intense Grand Challenge problems, NCEP's more unusual and stringent requirement is the need to generate operational forecasts according to rigid, predetermined schedules. Compounding the scheduling problem further, the skill or quality of a forecast is highly dependent upon the quality of the initial analysis. Accordingly, an operational model run is begun as late as possible in order to incorporate the largest possible quantity of observational data. The model is then run as rapidly as possible to meet its deadline.

The complex nature of NWP problems coupled with such operational requirements forces a sometimes uncomfortable accommodation. That is an acceptance that it is largely the marketplace that determines both when a new system should be obtained and what type of system it should be. It is the marketplace that sets the price/performance levels of commercially offered systems and that determines the most cost effective way to satisfy high-end computing needs. Weather modeling applications can be scaled up with relative ease. These changes of scale in NWP models have considerable practical advantage and will saturate any computational capacity that can be offered in the foreseeable future.

A number of topics that impact technology choices regarding high-end systems are discussed below.

Parallelizing the Applications

Since the advent of multiple CPU systems, NCEP has employed parallelism in the execution of its operational models. However, the granularity of this experience is not what is intended by the

use of the term "parallelism" today. True parallel processing is achieved at the array processing level and is therefore defined by the "DO LOOP" construct.

Since the introduction of the Cray YMP-8, NCEP has parallelized its operational models using all of the machine's CPU and memory resources to complete each operational code within the time dictated by product distribution schedules. The possible extent of this parallelization at the National Centers doubled with the introduction of the Cray C90 in 1993. Significantly greater improvements will result when MPP techniques can be applied to dynamic equations. These techniques will dramatically increase the degree of parallelism in model calculations.

Not all applications are easily decomposed for parallel processing. If the degree of parallelism is moderate or if the interaction among the pieces is extensive (in terms of synchrony, data sharing, and message passing), then the precepts of Amdahl's law prevail. In simple terms, this law states that no parallel process can be completed faster than the portion on the longest critical path.

The problem of converting from moderately parallel vector processing to massively parallel processing is one that has engrossed much of the high performance computing industry for the past decade. As a user of industry solutions, the NCEP has carefully followed these efforts and is presently running benchmark tests on some MPP systems. Still, there are risks in committing to this technology now even though it holds great promise for the future, perhaps even for the next NCEP supercomputer procurement. Continued participation in parallelization experiments is vital to the ability of the National Centers to make an informed assessment of the market.

Efficient Utilization of Large Numbers of Processing Nodes

The ability to effectively apply parallel processing technology depends largely on the ability to define a problem as a series of concurrent operations. This is true whether the implementation platform is moderately or massively parallel. Concomitant with parallelization is the need to keep all of the processing nodes busy. Idle processors are wasted resources. Moreover, the nature of large scale parallel processing is to create obstacles such as interprocessor wait times, message passing blocks, and data reference conflicts. Careful attention to performance optimization on MPP systems requires a significant investment in people with architecture-specific knowledge and with access to appropriate analytical tools.

Interprocessor Communication

March 13, 1995

A fundamental bottleneck on parallel processing systems is the need to synchronize operations and share data. This is a frequent stumbling block to the effective application of these systems. All parallel processing computers use multiple processors. Some use a shared memory model while others use local (distributed) memory. The communications bottleneck is most pronounced when systems are loosely coupled -- where each CPU is independent and autonomous -- but the effect is seen with all parallel system technologies. Each processing node contains resources to perform calculations independent of all other nodes until data or interdependent results are required. On many parallel systems the CPU, memory, and a connection to the outside world are self-contained with an internal network providing inter-nodal communications. An important task facing the application developer in this environment is to define the problem so as to minimize delays during synchronization. The greater the number of truly independent calculations that can be performed, the less conflict there is due to interprocessor communications.

Whether shared or distributed memory is chosen by the system designer, there are times when re-synchronization is required. Keeping system performance at an acceptable level often requires substantial effort to effectively distribute work over the processing nodes. Like classic traffic and routing applications, communication routes and transmission timings must be carefully planned to keep each node busy rather than standing idle awaiting new data.

Software

There are a number of key software issues relevant to high-end computing at the National Centers. The first of these is the operating system which must be a standard POSIX open architecture such as UNIX. Other important software concerns are discussed below.

Interoperability and Portability

An essential decision for any software is the computer language in which it is written. NCEP has standardized on ANSI compliant FORTRAN as its scientific programming language. As with most modern computer languages, the functionality of standard FORTRAN will continue to evolve and the language will increasingly offer modern programming constructs. This is particularly important in the case of high performance computing systems that utilize rather esoteric hardware features. To ease concerns regarding portability, vendors such as Cray Research sometimes use a meta-language that relies upon compiler pre-processing statements for a specific implementation of FORTRAN. With appropriate optimization instructions, such a pre-processor can interpret

March 13, 1995

these statements and use them to structure code specifically for the Cray architecture. Such pre-processing statements have the appearance of comments to other implementations of FORTRAN. Consequently the same source code can run at full optimization on the C90, for example, and still preserve integrity for execution on any other system.

Intelligent Compilers

The use of intelligent compilers that analyze the form of mathematical processes as represented by standard constructs has become commonplace. This type of analysis is important to the rapid deployment of NCEP operational codes. Each computer vendor customizes its code analysis during compilation in order to output executable code that performs best on a specific system. The interpretation of standard FORTRAN will be consistent among the offerings of all vendors, but the machine language instructions that are produced, as well as the commercial run-time service libraries, are highly customized to optimize performance. The performance of a specific class of applications such as the models that describe circulation patterns and motion within the atmosphere will differ widely among different vendor offerings. This is not only because of hardware differences, but also because some compilers do a better job of optimizing code for a particular class of applications.

Mathematical Subroutine Libraries

Another class of essential software is that of a collection of efficient, highly optimized subroutines that perform common, standard procedures. These are functions that offer well described techniques for solving a general class of problems. Typical of this genre is the Fast Fourier Transform (FFT), a common way of integrating motion fields in an array. The FFT is a fundamental mathematical tool that transforms gridpoint space into transform or waveform space to ease numeric computation. Because it is so fundamental, much effort has gone into optimizing it for various system architectures. There are whole classes of fundamental transformations used in the mathematical treatment of a particular class of problem and the challenge is to find a near optimal solution for a specific hardware configuration. MPP systems solve problems using an entirely different approach than that used on a vector processing system. The richness and speed of supporting libraries has a direct affect on the techniques used (since only those tools available for the architecture at hand can be employed) and the results obtained. Accordingly, for MPP systems to become competitive in the marketplace, vendors will have to provide a mature set of such routines.

Communications Infrastructure

March 13, 1995

One of the major issues at modern supercomputing centers is that of the use of very high speed inter-system communications facilities to link together the components of the center. Not only are these links used to connect supercomputers to file servers and mid-scale processing support systems, but they also connect host systems to network-attached disk arrays, message passing routers, near-line hierarchical storage media, and broadband video and workstation visualization systems. At the NCEP, interprocessor communications are undergoing a major revision. The revitalization of the NCCF in mid-1995 will include the replacement of two older Cray systems with two J916 models. Interlacing these new systems with the Cray C90 will, for the first time, allow experiments at the NCCF on distributing ensemble forecast models across multiple computer systems using high performance parallel interface channels. This will be possible, in part, due to the ability to share disks across all three Cray systems. This will permit an investigation into inter-system synchronization and into data sharing issues. It will also allow modelers to try new methods of processing and to investigate practical problems surrounding alternative computing architectures. Some of these, such as using parallel vector machines and ganging groups of processors into families to perform common distributed task processing, may offer advantages to future work at the NCCF.

Applications Software

In any large system, there are a number of software add-ons that extend the usefulness of the system. These range from interactive access packages such as X Windows to system administration and system performance packages. Interactive debuggers and special tuning software are essential to achieve maximum performance from applications software. Local data management and hierarchical storage management software are also required.

Traditional Mainframe Systems

Disposition of the HDS Systems

The overall strategy is to retire the HDS systems after moving their workload from the IBM-compatible MVS environment to that of an open architecture UNIX operating system. This migration will be phased over time to minimize operational disruption and to avoid having to convert codes that are already being rewritten as part of other NWS modernization programs. The migration will occur in four distinct phases:

March 13, 1995

1. Move the operational codes and data onto the Cray or other UNIX systems.
 1. Retire those portions of the HDS hardware and software that provide operational redundancy. Initially this will release the HDS EX-50 and its support equipment.
 2. Convert the remaining codes and processes from MVS to UNIX.
 3. Remove the remaining IBM-compatible components, including the HDS model EX-65.

The general schedule of these activities is as shown in the following table.

Task	Time Frame
Provide a suitable UNIX platform for migration activities	Late calendar year 1994
Move the operational processing	Mid-calendar year 1995
Remove the first HDS system (the EX-50)	Late calendar year 1995
Remove the final HDS system (the EX-65)	Late calendar year 1996

Historical Perspective

Although not strictly necessary to understand the changes that are in progress, an historical perspective can sometimes yield insight into strategic planning processes. This discussion briefly looks at how the National Centers arrived at this juncture and demonstrates why the HDS legacy systems can now be phased out.

The Hitachi Data Systems (HDS) line of IBM-compatible mainframe computers were installed at the NCCF beginning in 1983. These systems replaced three IBM model 360/195 computers that were acquired in 1973 as the computationally intense supercomputers of their time. The model 195 systems had been at the top of IBM's product line, and they combined general processing functionality with specialized hardware features that afforded high speed scientific computational capacity as well. Over this period of time, the NMC developed an integrated and unified processing capability built around the monolithic processing power of a single architecture.

In 1981, the model 195 systems were augmented by the acquisition of a vectorizing Class VI computer system, the Control Data Cyber 205. The fastest systems marketed at that time were produced by Cray Research and Control Data Corporation (CDC) and the Cyber 205 was selected through a competitive procurement. The Cyber 205 was a powerful numerical computer but it lacked the flexibil-

ity and robustness to serve as more than a special purpose vector processing engine.

There continued to be a need for a general purpose data processing system that could provide quality control and data management functions, as well as host an interactive programming test and development facility. The HDS systems were acquired to replace the IBM 360/195 systems but they were not intended to be used for the execution of numerical weather models. Rather, these systems performed a front-end supporting role for the specialized Cyber 205 computational systems where the predictive models ran.

The choice of HDS replacement systems evolved naturally from the IBM 360/195 systems and resulted from the computing economics that prevailed at that time. The support structure that had been developed over the preceding decade was preserved by selecting an IBM-compatible replacement system. There was an enormous range of software products compatible with the IBM architecture and IBM-compatible hardware was cost competitive. HDS processors and disks, STK tapes, disks and printers, as well as communications controllers and device switches were all compatible and easily intermixed at the NCCF.

The result of this mixed architecture was that the computational workload was split into two parts; a numerically intense and time critical portion, ideally suited for the supercomputing engine, plus a more conventional workload appropriate for a standard commercial architecture. This traditional workload benefited from an industry-standard architecture that could host a variety of common data manipulation and program management and development tasks. Moreover, a commercial system that had many offerings of off-the-shelf software and competitively priced hardware was important for maintaining concurrence with the computer industry. Many products were later added to the NCCF systems to improve productivity and enhance efficiency.

Over a period of years, the two initial HDS systems were enhanced using contract growth options and ultimately replaced via a series of interagency transfers of similar compatible systems and through contract engineering changes. The present configuration consists of one model EX-65 dual processor system and one model EX-50 single processor system. This is half of the former maximum configuration of three HDS 9000 systems plus the EX-50.

Declining role of the HDS systems

The National Centers are committed to employing a POSIX compliant Open Systems Architecture and using industry standard network communications. This commitment has already been instituted at

March 13, 1995

the NCCF in the form of the Cray C90, YMP-8 and EL systems. These systems provide the NCEP with a range of capacities, from supercomputer to file-server. However, due to the relative price/performance range of high-end supercomputers, they are less cost-effective when used as general purpose computers to perform non-numerical support functions such as certain kinds of quality control, data archiving and manipulation, statistical inference analysis, various pre- and post-processing tasks, and administrative functions.

In September, 1994 the NMC signed an engineering change to the Cray C90 contract. That change provides the Government with enhanced processing in the mid-scale computer range, reduced maintenance costs for the Cray processors and disk storage subsystems, and operating energy savings. Two Cray model J916 computer systems with new disk subsystems will be installed in mid-1995. These systems will provide an additional base for continuing the evolution from an IBM compatible platform to a UNIX-based POSIX standard architecture. When these systems are installed, the National Centers will have established an enhanced and cost-effective route for the conversion to UNIX.

There are significant hurdles to overcome as existing programming standards relative to data formats, processing management, and operational commitments are transferred from the MVS environment to open systems. Nevertheless, this process has already begun in the form of conversions for the C90 and YMP-8 systems as well as in the area of end-user training necessary to support these conversions. Planning is underway relative to the details of the migration of the HDS workload to UNIX so that this may be accomplished in a coherent and coordinated manner.

The importance of coordination in this regard cannot be overlooked. The NCCF is to some extent a service provider with clients and customers like any business. So far as it is possible to do so, clients should be encouraged to migrate through the availability of improved service and facilities. This is preferable to forcing a migration by decree. That approach would require users to reinvent already functioning software systems that are optimized for and ingrained into a specific operating system architecture. There are NWS-wide modernization programs that will obviate the need for some of the output produced by this interlaced collections of programs and it is therefore not economical for users to convert them at this time just to run on another platform. From the perspective of those users it is far better to reinvent their processes, rather than hastily patch constituent parts for a short-term move.

Approach to Migration

March 13, 1995

The approach to migration at the National Centers is to transfer the operational pre-processing and post-processing work performed on the HDS systems to UNIX as rapidly as possible, and to subsequently entice the remaining users to migrate using the following three techniques.

- The HDS users are dependent upon the output of the operational job suite. The products and services that they desire will become obtainable only (or most efficiently) from the UNIX target systems.
- The NCCF will provide UNIX platforms as targets for the migration; will make useful migration tools available; will assist in providing appropriate training; and will offer experienced guidance and assistance. In late 1994, the NCCF began a conversion pilot effort to move from the MVS environment to a UNIX variant for the HDS systems called Osiris. The purpose of this is to gain experience with a UNIX system designed to be operated on IBM-compatible processors. Data formats and conversion issues can be identified while remaining within the physical confines of the HDS hardware systems. This project is a continuing one, and is expected to be of benefit in identifying and resolving data conversion issues.
- The NCCF will gradually constrict the resources available to the legacy MVS systems. This effort will be managed by running a UNIX system on the same hardware processor (the HDS model EX-65) where the MVS system resides. As conversion progresses, the computational resources allocated for use by the MVS partition can be gradually transferred to the Osiris partition.

Past experience suggests that about seventy percent of the workload can be migrated quickly if there is interest on the part of the responsible program officials to do so. Much of the remaining thirty percent consists of legacy software that is well debugged and operational; and therefore in the payback phase of its system life. This work is valuable and contributes productively to mission objectives, but conversion efforts would be excessive, typically because other plans already call for the retirement of these programs or a discontinuation of the use of their output. This is particularly true with regard to the modernization efforts undertaken by the NWS, which call for new output data standards and formats.

Objectives of the Migration from MVS to an Open Architecture

March 13, 1995

The objectives of this migration effort are summarized by these points.

- Strive to provide a computing center that is an integrated processing facility rather than a collection of independent systems.
- Keep UNIX target systems open and flexible while gradually reducing the availability and performance of the MVS systems.
- Provide for an evolutionary migration by using a phased reduction of services.
- Accommodate gradual life-cycle changes.
- Provide for the graceful retirement of established, proven technology.
- Accommodate data-sharing and data-flow throughout the National Centers, with regard to both intersystem communications and external data services.
- Eliminate the redundancy that was necessary to support operational processing once the HDS systems cease to provide these operational functions. This allows for the early release of major components such as the HDS model EX-50, multiple device switching units, and channel coupling devices used for intersystem communications.
- Minimize the hardware and software logistic and maintenance support by eliminating complexities associated with shared disk systems and contingency operation provisions.
- Recognize and support the burden of legitimate legacy software comprised of fully functional operational systems that are in the payback part of their life-cycle.
- Provide a seamless and evolutionary migration rather than demand the conversion of mature, functioning software systems that are deeply intertwined with each other. These programs interrelate with compatible codes developed long ago, utilizing older data standards designed for legacy systems. Adaptation rather than the hasty conversion of mature codes is called for. Every platform has native data organization techniques and processing methods that allow codes to increase their efficiency when particular hardware and software features of its architecture are exploited.

What this approach offers is to allow the National Centers to:

- Gradually phase out older generation processing commitments
- Selectively continue support for declining technologies
- Provide a clear migration path, sharing proven conversion techniques and examples
- Encourage modernization by offering transition aids and by supplying adequate capacity for the conversion effort through the allocation of resources as they are needed

March 13, 1995

- Establish training programs developed from experience and to thereby offer assistance from path blazers
- Offer a well-defined phase-out date for legacy systems -- reaching this point via a gradual decline in support for the legacy software
- Offer the advantage of platform-independent data availability.

The National Centers are currently addressing many specific issues relevant to the migration to open systems. A detailed UNIX Conversion Plan is being prepared within the NCO. To assist in the preparation of this Plan, the NCO will take advantage of an agreement with the Software Engineering Institute of Carnegie Mellon University. This agreement, available through a program sponsored by the System Engineering Staff of the NOAA Systems Program Office, provides products and services to facilitate the transfer of software engineering technologies to NOAA. An initial descriptive document that details some aspects of the conversion problem is included as Attachment 8. An initial, simplified version of the finished plan may be ready by the spring of 1995 and will provide schedules and milestones for the systematic conversion to an open UNIX computing environment.

Networking and Communications

There are three fundamental components of the networking and communications architecture of the National Centers:

- Communications among the National Centers (WANs)
- Communications within a National Center (LANs)
- Communications with organizations external to the National Centers.

Each of these requires separate consideration and is addressed below.

Internal Communications -- Wide Area Networks

A fundamental assumption of the reorganization of the National Centers is that communications between the Centers will be transparent. That is, as it regards their routine work it will not matter to the employees at any Center whether or not they are geographically near the people, computer systems, or data sources of any particular Center. A reliable WAN with sufficient bandwidth is required to make this assumption a reality.

The planned topology is a fully redundant network in which the Centers are interconnected using dedicated leased T1 lines. This design ensures that communications among the Centers will not be interrupted by a failure at any one site. Each Center will have the ability to communicate independently (that is, without dependency on an intermediary) with any other Center. The details of this WAN topology, along with an evaluation of alternatives, can be found in Attachment 5.

Internal Communications -- Local Area Networks

At each NCEP site the computing systems will be connected via LANs. Each site will employ a high bandwidth backbone, intelligent hubs, and routers to support local users, processes, and equipment. As in the case of wide area networking, each LAN will rely on TCP/IP ethernet communications.

Current LAN technology at the National Centers uses fiber optic cable and the fiber-distributed data interface (FDDI) in a ring topology to provide 100 megabit per second backbone capacity. However, current technology demonstrates that comparable data rates can be achieved over twisted pair copper wire using CDDI (copper-distributed data interface); and that much higher rates can be attained over fiber optic cable using ATM. Since FDDI, CDDI, and ATM all support TCP/IP ethernet, any of these can and will be used. Specific choices in specific cases will be made based upon bandwidth requirements, existing cabling systems, and cost.

Within a given site, LAN segments are connected via intelligent network hubs. A LAN segment typically supports a workgroup determined by physical or organizational proximity. The hubs perform routing and bridging functions and are designed and configured to optimize network performance and minimize contention. The manner in which an individual device is connected to a LAN segment depends upon the individual communications requirements of the processes resident on that device. Standard ethernet, FDDI, CDDI, and ATM are all consistent with this network topology. Finally, each LAN is linked to the WAN via network routers.

External Communications

Operational communications between the National Centers and external customers are managed primarily by the Office of System Operations (OSO). Currently, communications between NCEP Central Operations and OSO is by way of legacy MVS mainframe systems utilizing the Systems Network Architecture (SNA) protocol. The

March 13, 1995

National Centers would prefer to modernize that link and employ standard, open TCP/IP ethernet, making the communications independent of any specific computer or vendor.

OSO operates the Automated Field Operations System (AFOS) and the National Centers have access to AFOS, including access to the communications subsystem that it relies upon. The AFOS communications protocol is ADCCP (Advanced Data Communication Control Procedures). The National Centers only employ ADCCP incidentally when using AFOS. When the Advanced Weather Interactive Processing System (AWIPS) replaces AFOS, then ADCCP will be replaced by TCP/IP ethernet, at least with regard to internal AWIPS LAN communications. This will make AWIPS much more compatible with other NCEP systems than AFOS has been.

Another communications protocol used by the National Weather Service (principally by OSO) is X.25. This international standard is widely used by WMO members, including OSO, to support data requirements in real-time. To obtain real-time data from OSO, NCEP must conform to OSO's procedures. The National Centers will use X.25 and dedicated low-speed telephone lines for this purpose. Such X.25 communications links will terminate in UNIX workstations which will ingest the data and make it available to users on the LAN via TCP/IP ethernet.

The National Centers also communicate directly with a small number of users such as the Federal Aviation Administration. These communication links represent legacy systems, sometimes using non-standard protocols and services, and it is intended that they will be upgraded gradually to TCP/IP ethernet.

Non-operational external communications involving the NCEP will use standard TCP/IP ethernet. In those cases where this is not yet the case, modernization efforts will move in this direction over time.

Data Flow and Data Handling

The data system at the NMC was designed to support the operational use of numerical weather prediction models. Accordingly, it evolved to meet the schedule required for that purpose. The mixture of capabilities in place today are a result of either the need to augment this system or the attempt to force this system, designed for a specific purpose, to meet a wide range of demands as represented by many different groups within the National Centers.

March 13, 1995

Operational forecasters at each National Center need access to data. In the case of observational data this is a real-time requirement, unlike the need for data to support numerical models. At the NMC in Maryland, because of the collocation of central computers, model developers, and forecasters, the approach has traditionally been to make the data flow that was designed for modeling also suffice for the forecasters. That has never been very satisfactory. At NSSFC and NHC alternate, independent paths and systems have been developed for real-time data. Forecasters at each National Center also have access to AFOS as a source for data.

In the cases of model output, satellite imagery, derived graphics, and so forth, a variety of methods for delivering data to the National Centers has evolved over time. These methods depend upon a diverse assemblage of equipment, data formats, and software. Managing and operating this heterogeneous collection of capabilities is time consuming and labor intensive.

The National Centers are presently designing a comprehensive data flow and data handling system. This system will be consistent with the overall standards defined for use at the National Centers. The key points relevant to this data system are these:

- Data formats -- gridded data will be stored in GRIB and point data will be stored in BUFR.
- All processing -- decoders, pre-processing, and post-processing -- relevant to the data will be performed on UNIX machines.
- The database will be maintained on UNIX machines utilizing standard UNIX file structures.
- The National Centers will continue to monitor developments regarding commercial database management systems, although there are no plans at this time to incorporate such products into the operational NCEP data system.
- Communications within and among the National Centers will rely upon TCP/IP ethernet.
- Operational communications between NCO and OSO is currently based upon SNA, a proprietary protocol that supports direct communications between mainframe systems. This data exchange process will move to standard TCP/IP and UNIX as both NCEP and OSO are able to move operations away from their legacy MVS systems.
- The National Centers will use the data communications systems that evolve to support AWIPS when possible.
- The modern NCEP data system will incorporate a variety of interactive, user-friendly, graphics-oriented tools to enable users to access, display, and manipulate data. These tools will be part of N-AWIPS.

March 13, 1995

- The same data system will serve each National Center.

The NCEP data system will serve the needs of the scientists and forecasters within the National Centers. It should also suffice for the needs of others who are granted access to work as clients within the NCEP computing environment. External users will be given limited access to a subset of the NCEP data sets via ethernet servers without the data itself being reformatted for those users. Customers for whom it is necessary that NCEP tailor data will generally receive data products rather than the data itself.

Scientific Workstations

Scientific workstations at the National Centers are used primarily to view graphical representations of data and information, including the output of numerical weather prediction models. The ability to present complex information to a user graphically and to enable that user to interact with that information in a wide variety of ways is of fundamental and growing importance. The sheer volume of information and the variety of ways that analysts, forecasters, researchers, and programmers must interact with it require the development of electronic tools for these purposes.

Workstations provide users with a host of specific and interrelated capabilities. Using windowing techniques a user can apply several capabilities simultaneously so that work can be done on several tasks in parallel at his or her pace. This is far more productive than work done serially at the system's pace as in the traditional mainframe computing environment. The productivity gain made possible by a modern scientific workstation environment over that of a mainframe system is tremendous. For example, using a scientific workstation a developer can edit source code in one window; compile, debug, and execute it in another; monitor the use of system resources in a third; and view a graphical presentation of the results in a fourth. Windows can also be maintained to support a standard suite of personal productivity tools like e-mail and word processing, to access services from remote systems, and to initiate new tasks while waiting for various other processes to complete. Furthermore, these tasks can directly and immediately utilize resources anywhere on the network. The user is limited only by his or her ability to keep track of these activities and not by the systems themselves. This rich environment supports rapid prototyping and end-user

computing because the user can take advantage of a very wide range of capabilities and can quickly see results.

By insuring that the methods employed with regard to workstations are standard and open, scientists at the National Centers can easily collaborate with others at universities and at diverse research laboratories. Models from other centers such as the European Center for Medium-range Weather Forecasting (ECMWF) can be used, and their output viewed, as conveniently as can NCEP models. The commonality and openness of this environment produces a synergy that fosters accelerated development through the broad sharing of data, ideas, and results.

Production Control in a Distributed Environment

The System Monitor and Scheduler (SMS) is a software facility that was developed by the ECMWF to provide centralized scheduling and monitoring of complex network job streams. This is an X/Motif application, written in C, and designed to run in an open network that includes supercomputers, mainframes, servers, and workstations. The National Centers are evaluating SMS for production job control and event scheduling. This tool, or something much like it, will be implemented in 1995 to help control operational jobs with the NCCF.

SMS provides a graphical user interface by which computer operators and others can:

- schedule and control operational jobs such as the execution of NWP models
- monitor the routine progress and status of those jobs
- intervene to correct, rerun, or resequence jobs in the event of abnormal conditions.

Since the SMS works in an open network environment, this tool can be used for interdependent programs running on a number of designated systems. As operations extend over many network nodes, all scheduled and monitored by SMS, it will become relatively easy to direct a specific application to run on the platform most appropriate to it, whether that is a supercomputer or a small UNIX server. This will simplify the allocation of work to various computer resources and help to optimize the effective use of NCEP systems.

AWIPS

A key component of the National Weather Service's MAR program is the Advanced Weather Interactive Processing System which will be a highly automated and integrated weather information processing, communications, and display system. AWIPS will be deployed at each Weather Forecast Office, at each River Forecast Center, and at each National Center.

AWIPS at the National Centers requires a number of capabilities beyond those required by a WFO or RFC. A brief inspection of NCEP operations reveals that the forecast processes at the national level require information spanning extended spatial and temporal scales. In addition, output for experimental as well as operational NWP models must be available at the National Centers for the full temporal and geographic model extent. The differences between the AWIPS requirements at a field office and those at a National Center are recognized explicitly in the AWIPS contract. The National Centers are responsible for the development of the National Centers unique capabilities. Furthermore, the National Centers are responsible for the integration of those capabilities with the national AWIPS system.

The NCEP approach is to develop the National Centers AWIPS (N-AWIPS) system through a series of prototypes. The Pilot Operational Projects (POPS) approach represents a very aggressive effort to quickly provide a limited subset of N-AWIPS capabilities to forecasters at the National Centers. These capabilities will be iteratively extended and then integrated with AWIPS as the national system is delivered and upgraded by the prime AWIPS contractor. N-AWIPS capabilities will also be distributed throughout the National Centers and elsewhere for use on non-AWIPS equipment. This equipment encompasses a wide range of scientific workstations including those acquired under this information technology plan. N-AWIPS will be supported not only for analysts and forecasters, but for a wide range of non-operational users, including modelers and programmers.

N-AWIPS is intended to extend the capabilities currently provided at the National Centers by a variety of largely unrelated systems. The primary components of these existing systems are:

- VAS Data Utilization Center (VDUC) for display of graphics and imagery, data ingest, and generation of diagnostics
- Manual Graphics System (MGS) for data display, manual graphics generation, and product launching

- AFOS for access to data
- NMC Plotting System for hard-copy plots of data and model output.

As N-AWIPS is deployed and becomes operational, all of these systems will be phased out at the National Centers.

Contract Support

The National Centers routinely augment staff by employing contract support in two distinct manners. First, contractors are used to perform specific tasks requiring expertise in narrow domains that are generally outside the area of interest and experience of NCEP employees. For example, when the National Centers acquire a new supercomputer, its manufacturer typically provides a small number of systems and applications programmers with extensive experience in the architecture of that system. These contractors assist NCEP scientists and programmers in migrating applications to the new environment and in structuring codes and operational job suites so as to best utilize computing resources. The period of performance for contractors of this type is typically a few to several years and roughly coincides with the life of the system they support.

A second type of contract support is of a more general nature. Most of the Government scientists at the National Centers use computing systems to perform a range of tasks related to specific programmatic areas. These areas are highly correlated with individual backgrounds and expertise. The employees involved range from forecasters who use computer systems to make operational forecasts to researchers who work to improve the performance of numerical weather prediction models. For the majority of these people, the NCEP computing and networking systems are simply tools that support their efforts to do specific jobs. By contrast, some employees, most notably within the NCO, focus their attention directly on the computing infrastructure. Within this environment there are some support activities of a general but transitory nature such that it is cost effective to provide that support through contractors. An example of this sort of activity is that of software support for the development of AWIPS at the National Centers. The N-AWIPS system will be built over a period of several years by the National Centers. This effort will require as many as twenty extra people for system development and implementation. Employees of the National Centers will direct the development and participate in it, but most of the work can be performed by people with general backgrounds in

March 13, 1995

systems analysis, programming, system administration, and so forth. These are skills that are readily available from many commercial sources. Over a period of approximately 5 years a team of developers, mostly contractors, will build and deploy N-AWIPS. After that time the Government can terminate the N-AWIPS support contract and maintain the system with a much smaller staff.

The use of contractors to work on tasks of a general nature to institute specific NCEP programs typifies this second type of contract support.

The management of the NCEP anticipates that at any time there will be a small number of contracts in place to support a variety of activities at the National Centers. The precise nature of the contract support will change over time but the need for support of the two types described above will continue indefinitely.

External Coordination

The computing architecture of the National Centers, as described in this document, will change substantially over the next few years. These changes will have a direct impact on a number of organizations and individuals external to the National Centers. Some of those most affected will be NWS components such as the Office of System Operations and the Office of Hydrology; as well as other NOAA line offices, particularly NESDIS. In addition, the National Centers work with components of the Department of Defense to provide mutual support. Substantive changes in the systems of the National Centers must be coordinated with all of these.

Employees of the National Centers routinely discuss issues such as those that make up this plan with their counterparts in other organizations. In addition to those informal conversations, this document will be disseminated to other NWS units as well as to selected NOAA line offices upon its approval. During the second half of 1995, representatives of the National Centers will meet with the staffs of interested external organizations to discuss the schedule and impact of planned technological changes. Thereafter, periodic meetings should be held to keep all interested parties apprised of the direction and rate of change; to gauge the impact of change; and to ensure that essential or expected services are not unexpectedly disrupted. The National Centers will assume responsibility for arranging regular meetings of this sort.

March 13, 1995

3. RELATIONSHIP TO THE NOAA STRATEGIC PLAN

The current NOAA Strategic Plan identifies three key program elements that are directly supported by the NCEP through its computing and communications systems. NOAA's strategy is to create an integrated environmental observation, assessment, and forecast service that supports the Nation's economic and environmental agenda by:

- Significantly improving short-term (immediate to 60 days) forecast and warning products to protect the American public;
- Implementing reliable seasonal to interannual (60 days to 10 years) climate forecasts to guide business and economic decisions;
- Developing science-based policy advice on decadal to centennial (10 years to 100 years) changes in the global environment as a basis for National policy decisions on economic security.

Advance Short-Term Warning and Forecast Services

Advanced supercomputers are needed to process the data from new observing systems such as Doppler radars, automated surface observing systems, and advanced satellites. They are also needed to refine the temporal and spatial resolution of the numerical models that utilize the data as well as to support improved model physics. The planned processing architecture of the National Centers is designed to provide computing services consistent with this key NOAA strategic element.

The reorganization of the National Centers will result in centers that are specialized to deal with particular short-term forecast problems. These include hurricane, severe storm, aviation, and marine forecasting. In order to effectively support these geographically dispersed centers, each relying on numerical models run at a central location as well as upon input from each other, it is necessary to link them via a modern high-speed communications infrastructure. That infrastructure is described in this document and is consistent with the NOAA Strategic Plan.

Modern workstations are essential in order to allow forecasters to use a wide range of environmental data to full advantage. This document describes the open systems approach employed by the National Centers to develop meteorological applications on such

March 13, 1995

workstations in a way that is consistent with the approaches taken both in the national AWIPS program and in the larger community of environmental science.

Implement Season to Interannual Climate Forecasts

NOAA plans to implement predictions of seasonal to interannual variability of the coupled ocean-atmosphere climate system. The forecasts provide information on temperature, rainfall, and ocean circulation patterns of substantial benefit to the agricultural, energy, fishing, and other sectors. (Quoted from the 1994 NOAA Strategic Plan.)

The initial target of this key NOAA strategic element is ENSO prediction. The large-scale computational systems of the National Centers provide a crucial resource necessary to achieve this goal. Coupled ocean-atmosphere models now run routinely at the National Centers and these form the basis of current operational multiseason forecasts as described earlier in this document. As in the case of short-term warning and forecast services, several of the components of this plan, including the described organization of the NCEP itself, combine to help meet this goal. These include the communications infrastructure, the scientific workstations, and the methodologies employed by the National Centers.

Predict and Assess Decadal-to-Centennial Change

An essential ingredient for NOAA to realize progress in this strategic element is a consistent analysis of the existing meteorological record. Data analyses from the past exhibit apparent anomalies that result from changes unrelated to the data itself. Specifically, as understanding of climate has improved and as the tools that are used to analyze relevant data have changed, so too have the resultant pictures that portray climate conditions. The National Centers are participating in a continuous reanalysis of the historic record using modern numerical models and computers to recast that record in the light of current understanding and capabilities. This effort will remove apparent anomalies from the climate record, making real changes in the climate more amenable to detection and understanding. This is a widely cooperative effort involving resources from several nations, a number of governmental agencies, and many private organizations.

March 13, 1995

4. ALTERNATIVES CONSIDERED

The overall NCEP system architecture today embraces a host of computing technologies -- supercomputers, traditional mainframes, network servers, scientific workstations, and personal computers. It also encompasses a wide range of networking approaches and protocols -- SNA, X.25, TCP/IP ethernet, FDDI, Internet, ATM, LANs and WANs, leased lines, microwave, and dial-up telephone services. The fundamental question is how to simplify and standardize this environment in order to better support it and make the most effective use of resources.

The needs of the National Centers are extensive and various. A great many alternatives have been and will continue to be considered and used to meet those needs. Experience has shown that several alternatives can be simultaneously, effectively, and appropriately matched against specific needs to provide necessary services provided that those alternatives share certain characteristics.

In order for a diversity of systems to be manageable, the National Centers have evolved a number of standards over the last several years. Adherence to these standards has helped to guide or direct the selection from among alternatives in many specific cases. It has also helped to identify practices and methodologies that should be changed in order to achieve consistency throughout the systems supporting the National Centers. These standards, which have already been described in some detail, are:

- Open computing and network architectures
- TCP/IP ethernet protocol
- NFS and NIS networking services
- The UNIX operating system
- FORTRAN and C programming languages
- X Windows and Motif for graphical user interfaces
- GRIB and BUFR data formats.

After stating these concepts and establishing these standards, there is one particular area in which there are no true alternatives. Weather forecasting is one of the historic grand problems of computing and no foreseeable system will be large enough and fast enough to solve it. There are no alternatives in this regard but to work with industry to utilize the most powerful systems available near the forefront of supercomputer technology. It is this path that the National Centers will continue to follow.

March 13, 1995

In all other areas, specific alternatives will be identified and considered when individual procurement actions are undertaken. No attempt is made here to identify them.

5. SECURITY

Attachments 6 and 7 are the security plans for the NOAA Central Computer Facility and for the NOAA Science Center respectively.

6. ACCESSIBILITY

There are no particular requirements for special accessibility but special accommodations can be made when necessary.

Acronyms and Abbreviations

ADCCP	Advanced Data Communication Control Procedures
AFOS	Automated Field Operations System
ANSI	American National Standards Institute
ASOS	Automated Surface Observing System
ATM	Asynchronous Transfer Mode
AWC	Aviation Weather Center
AWIPS	Advanced Weather Interactive Processing System
AVN	Aviation Global Model
BUFR	Binary Universal Form for the Representation of Meteorological Data
CDAS	Climate Data Assimilation System
CDDI	Copper-Distributed Data Interface
CDC	Control Data Corporation
COAM	Coupled Ocean-Atmosphere Model
CODAPS	Coastal Ocean Data Assimilation and Prediction System
COTS	Commercial, Off-the-Shelf Software
CPC	Climate Prediction Center
CPU	Central Processing Unit
ECMWF	European Center for Medium-range Weather Forecasting
EDAS	Mesoscale Eta Data Assimilation System
EMC	Environmental Modeling Center
ENSO	El Nino/Southern Oscillation
ERL	Environmental Research Laboratories
ETA	Early Eta Forecast Model
FDDI	Fiber-Distributed Data Interface
FFT	Fast Fourier Transform
GDAS	Global Data Assimilation System
GFDL	Geophysical Fluid Dynamics Laboratory
GMMM	GFDL Multiply-nested Movable Mesh
GOES	Geostationary Operational Environmental Satellite
GRIB	Gridded Binary
GSFC	Goddard Space Flight Center
GUI	Graphical User Interface
HDS	Hitachi Data Systems
HPC	Hydrometeorological Prediction Center
IBM	International Business Machines
IP	Internet Protocol
LAN	Local Area Network
MAR	Modernization and Associated Restructuring
Meso-ETA	Mesoscale ETA forecast model
MGS	Manual Graphics System
MOD	Meteorological Operations Division
MPC	Marine Prediction Center
MPP	Massively Parallel Processing
MRF	Medium Range Forecast Model

March 13, 1995

MVS	Multiple Virtual Storage
N-AWIPS	National Centers AWIPS
NASA	National Aeronautical and Space Administration
NCCF	NOAA Central Computer Facility
NCEP	National Centers for Environmental Prediction
NCO	NCEP Central Operations
NDA	Network Disk Array
NESDIS	National Environmental Satellite Data and Information Services
NEXRAD	Next Generation Weather Radar
NFS	Network Files System
NGM	Nested Grid Model
NHC	National Hurricane Center
NIS	Network Information Services
NMC	National Meteorological Center
NMFS	National Marine Fisheries Service
NOAA	National Oceanic and Atmospheric Administration
NOS	National Ocean Service
NOW	NOAA Ocean Wave
NSSFC	National Severe Storms Forecast Center
NWP	Numerical Weather Prediction
NWS	National Weather Service
OAR	Office of Atmospheric Research
OI	Optimal Interpolation
OSO	Office of System Operations
POSIX	Portable Operating System Information Exchange
QLM	Quasi-Lagrangian Model
RSM	Regional Spectral Model
RUC	Rapid Update Cycle
SEC	Space Environment Center
SMS	System Monitor and Scheduler
SNA	Systems Network Architecture
SPC	Storm Prediction Center
STK	Storage Technology Corporation
TCP	Transmission Control Protocol
TPC	Tropical Prediction Center
WAM	Wave Model
WAN	Wide Area Network
WFO	Weather Forecast Office
WMO	World Meteorological Organization

Attachment 1

Statement of Intent between the National Oceanic and Atmospheric Administration and the National Aeronautics and Space Administration Regarding the Development of a NOAA Operations and Research Center at the NASA Goddard Space Flight Center

Attachment 2

Advances in Operational Numerical Weather Prediction: An Outlook for the Early Twenty-First Century

Attachment 3

An Overview of Operational Prediction Capabilities at Fleet Numerical Meteorology and Oceanography Center

Attachment 4

Memorandum Regarding Massively Parallel Processing at NCEP

Attachment 5

NCEP Wide Area Network Implementation Plan

Attachment 6

Security Plan for NOAA Central Computing Facility General Support Systems

Attachment 7

Security Plan for the NMC Workstation Facility

Attachment 8

NCEP UNIX Conversion Plan -- DRAFT

March 13, 1995